Specifying EO satellite campaign requests to meet science goals

Jennifer Dungan[a] and Robert Morris[b]

[a] Earth Science Division, MS 242-4, NASA Ames Research Center, Moffett Field, CA 94035 U.S.A., Jennifer.L.Dungan@nasa.gov;

[b] Computer Science Division, MS 242-4, NASA Ames Research Center, Moffett Field, CA 94035 U.S.A.

ABSTRACT

In the past decade, the number of Earth observation satellites has burgeoned. EO missions have been conceived to study different aspects and interacting pieces of the Earth system and scientists are designing increasingly complex, interdisciplinary campaigns to exploit multiple EO assets. In general, each EO mission has its own observation scheduling system, leading to a diverse array of specialized scheduling systems each tailored to a mission's unique requirements and characteristics. There is therefore the opportunity to allow planning and scheduling algorithms to assist a scientist requesting a set of observations, optimized over multiple sensors, to meet a science goal. This paper addresses the formal specification of a request to allow such algorithms to be exploited.

The most basic parameters of a campaign request include the list of sensors required or desired in combination with the time frames or intervals and geographic regions for which the observation(s) are needed. Temporal and geographic parameters may be dependent on exogenous events, such as the occurrence of a fire, hurricane, volcanic eruption, or seasonal event. The scientist would also like to specify constraints on the needed quality of observations (e.g. cloud cover) and may have dependencies between parameters. Fulfilling preferences (e.g. new data are preferable to older data) may improve the quality of subsequent scientific analysis. The cost of the set of observations will also be an important constraint. In this paper, we propose a structure for expressing these parameters and constraints. We use a fire emissions model validation campaign as illustration.

Keywords: Earth Observation System, Earth science goals, remote sensing campaigns

1. INTRODUCTION

There are currently scores of Earth observation satellites orbiting our planet. Some of these satellites have single sensors on board, others contain a suite of sensors using different technologies for passive and/or active remote sensing. Many satellites are government-sponsored, including the fifteen that are part of NASA's Earth Observing System (http://eospso.gsfc.nasa.gov); others are commercial ventures with data available for sale.[1] For scientists or science teams engaged in complex, interdisciplinary experiments utilizing data from many sensors, scheduling future observations can be a nontrivial exercise involving the interface with several separate mission teams and/or observation request systems.

In July, 2003, the first Earth Observation Summit took place in which thirty-three nations and the European Commission agreed to create the means to develop comprehensive, coordinated and sustained Earth observation. Representatives at the summit affirmed that timely, high quality, long-term, and global information is necessary for sound decision making. Several important goals were recognized: to monitor the state of the Earth continuously, to improve understanding of dynamic Earth processes, to enhance prediction of the Earth system and to implement environmental treaty obligations. A ten-year plan is being developed and is slated to be unveiled at the third summit in February of 2005. The framework describing that plan[2] states that the Global Earth Observation System of Systems (GEOSS) should allow "coordination and facilitation of the development and exchange of observations and products between members and relevant international and regional organizations." To make coordination possible, tools will need to be developed to facilitate the acquisition of data from multiple sensors on multiple satellites to satisfy goal-oriented requests.

The purpose of this paper is to describe the general concept of a remote sensing campaign and a structure for specifying observations to be requested across a constellation of Earth observation satellites. This paper also describes an architecture for a software system that enables a community of Earth Scientists to submit campaigns to a collection of sensors. The architecture provides the basis for planning technology that performs optimized campaign planning.[3,4] There are several ways this technology could be deployed – over a limited range of satellites under the control of a single space agency, as part of an infrastructure of an "intelligent" sensorweb[5] or as a component in the GEOSS.

We use the term "*campaign*" to refer to a systematic set of remote sensing activities undertaken to meet a particular science goal. A few examples of campaigns funded by NASA to support broad science goals are:

- the Multi-sensor Airborne Campaign Hydro, conducted over the Mahantango watershed in 1990 in Pennsylvania to study the remote measurement of soil moisture,
- the Pacific Exploratory Mission in the western Pacific (PEM-West), conducted in 1991 and 1994 to investigate the atmospheric chemistry of ozone production and fate and marine versus terrestrial sources of sulfur,
- the Boreal Ecosystem Atmosphere Study (BOREAS), conducted in 1993 and 1994 at two sites in northern and southern central Canada to study exchanges of radiative energy, sensible heat, water, CO_2 and trace gases between the boreal forest and the lower atmosphere, and
- the Large-scale Biosphere Atmosphere Experiment in Amazonia (LBA), begun in 1998 to study how Amazonia functions as a regional entity within the larger Earth system and how changes in land use and climate will affect biological, physical, and chemical functioning of the region's ecosystem.

There exist scores of others, including formal, agency-sanctioned or well-recognized efforts and many others that involve individual scientists or small teams attempting to answer scientific questions. All of the above-mentioned campaigns involved data from aircraft (suborbital) sensors and ground observations in addition to observations from satellite sensors. Aircraft and ground observations are usually more complicated to model than those from satellites because they represent unique configurations of research instruments and are often never repeated. For example, the ACE-Asia campaign[6] involved many months of schedule negotiation involving tradeoffs among aircraft engineering, travel and weather constraints. We confine our discussion here to satellite assets.

2. SCIENCE GOALS

Science goals are stated with respect to the geophysical or biophysical quantities (often called parameters or variables) that are to be measured. The twenty-four parameters that are the goals of the Earth Observing System set out by NASA,[7] include land surface temperature, vegetation, ice and cloud cover, energy flux, atmospheric concentrations of particles and gasses. Each of these parameters is actually a complex of individual quantities – for vegetation for example there are products representing leaf area index, the fraction of photosynthetically active radiation, net primary production and fraction of tree cover quantities. These specific quantities are not directly observable using satellite-borne sensors, but are inferred using sensor data on reflected radiation in various spectral bands and models of the processes. Therefore, goal specification in terms of quantities implies that processing algorithms exist to produce them from satellite sensor data.

Many science goals are directed toward the study of ongoing Earth processes, such as seasonal ice melt, vegetation green-up and senescence, agricultural yield dynamics, the spread of invasive species and air pollution transport. Toward these goals, campaigns may be constructed to assemble regular observations timed to capture snapshots of targeted regions when change is occuring. Other science goals are based on Earth system phenomena that are transient, ephemeral, unpredicted or simply unknown. Flood, drought, fire, volcanic activity, algal blooms, storms, insect defoliations, and oil spills among many others can be considered "uncontrollable" events. The time of their occurrence is not known beforehand, so the *deterministic* scheduling of observations is not possible. Goals may involve both ongoing and transient phenomena (e.g. to understand the effect of drought on fire, the effect of deforestation on flood magnitude or the the effects of dust transport and deposition on algal blooms).

3. CAMPAIGN SPECIFICATION

To generate campaigns to meet science goals, the minimum set of constraints needed to specify one observation includes the sensor to be employed and the timing and locations of its observations. The elements of a campaign specification structure therefore include these three items as well as specification of the exogenous events or processes that are the subjects of the science goal(s). The latter is under the control of the requester – the scientist(s) generating the campaign requests – and may not be explicitly stated in a formal request for a scheduled activity.

3.1. Sensor type

In the current era of Earth observation, most spaceborne sensors are unique, relatively fixed assets with few configurable components. Therefore, to specify the type of an observation, it is sufficient to give the name of the sensor. Exceptions to this rule include the Moderate Resolution Imaging Spectroradiometer (MODIS); since it exists on both the Terra and the Aqua satellites, its platform must be specified. Requests to the Tropospheric Emission Spectrometer (TES) on the Aura satellite must specify limb or nadir-observing modes. Lists of current Earth observation sensors can be found in several recent catalogs.[8–10]

3.2. Temporal constraints

Temporal constraints define restrictions on when an event is to occur, either in absolute terms or in relation to other events. Typically, temporal constraints are defined as a time interval. Because Earth-observing sensors image the region at nadir or within fairly narrow boundaries off-nadir, there is no flexibility to observe a given area at more than one time within a day. Therefore, to the extent that a given sensor has a fixed Earth-viewing schedule, it is sufficient to specify a time interval in terms of days.

3.3. Geographic constraints

A goal implies some study area which may be composed of one large region or multiple, disjoint regions. Any observation to be requested within the study area may have its own unique geographic specifications. Regions may be somewhat vaguely defined, such as the region centered on a point, or may have precisely specified boundaries. To specify regions, center points, simple geometric figures (circles or rectangles) or arbitrarily complex polygons represented at some cartographic scale could be chosen. Commonly, bounding boxes are the geometric unit of choice used to define regions because of their simplicity in user interfaces though they are sometimes not specific enough to preclude the acquisition of data outside the true area-of-interest. For a campaign involving more than one observation, multiple points, bounding boxes or polygons must be able to be specified.

3.4. Other constraints

In addition to temporal and geographic limits for each sensor's data, the user may have requirements related to the qualities of the image, such as the lack of cloud cover. There may also may be limitations on the overall cost of the total set of observations for those cases involving commercially available observations.

Finally, it is common to draw a distinction between observations that will be critical to meeting a science goal, and others will be desirable but will not cause the campaign to fail if they are not taken. A similar distinction may be drawn between constraints. For example, a user may request a specific sensor for an observation, but accept other sensors if the requested one is not available.

3.5. A campaign scenario

To illustrate the structure of a campaign, we present a hypothetical campaign based on the goal to test an emissions model predicting the aerosols released by wildfires. The broad outlines of this campaign were inspired by the Fire Locating and Modeling of Burning Emissions project.[11] Let us say the location of this campaign is in the southern California region, San Diego County. Data on several variables must be gathered in order to accomplish the analysis. In particular, vegetation type or biomass, atmospheric aerosol concentration and burned area are needed for the region. Fuel moisture content is a variable that also would be useful for the objectives of the science, though not a necessity.

There are several sensors that provide products at various spatial resolutions relevant to these variables. Landsat Enhanced Thematic Mapper+ (ETM+) or Thematic Mapper (TM) can be used for mapping vegetation type. Optimal timing for acquiring Landsat data for this purpose in Southern California would be the prior June or July in the same year that the fire burned, when forested land can most easily be spectrally distinguished from grassland. For mapping aerosol concentration, images coincident to burning must be obtained. MODIS on the Terra and/or the Aqua satellites would provide data for this variable. MODIS data from either platform could also be used to provide coarse spatial resolution burned area after (though not too long after) the fire were out. Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) or Landsat TM data would be desirable for mapping burned area with fine spatial resolution. For mapping vegetation moisture content, hyperspectral data from the EO-1 Hyperion instrument are relevant. The most useful data for this purpose would need to be acquired just preceding the fire.

3.6. From campaigns to plans

Informally, a "plan" is a set or sequence of activities that collectively achieves some goal. It is natural to think of campaigns as plans for accomplishing an Earth science goal, where each activity in the campaign consists of an observation. Recent developments in automated systems that perform planning have led to the notion of a "flexible plan", a plan where activities and their relationships are expressed as constraints, resulting in choice on the how the activities in the plan are to be executed. Flexible plans are motivated by the need for robustness in plan execution − if there is uncertainty in the environment within which the plan is to be executed, there is a need for flexibility in how plan activities are to be performed. For an example in the Earth observation domain, if sensors tend to be oversubscribed (i.e. more requests for time are submitted than can be serviced) then there should be flexibility in which sensor is used to do an observation in a campaign. Because we assume that execution environments for campaigns display such uncertainty, it is useful to associate campaigns with flexible plans.

Flexible plans are commonly represented at networks, where nodes stand for start times or end times of events, and there are edges between them representing temporal constraints. Figure 1 is a flexible plan representation of the fire campaign. Each node in the network represents an activity with an associated sensor type and geographic regions. Arrows represent the ordering of activities, with brackets representing temporal constraints in units of days. The earliest activity is by definition the campaign specification, in this case assumed to occur on November 15. Vegetation observations should be scheduled between 75 and 135 days after this initialization date. The start of the fire will be estimated to occur between 80 and 350 days after initialization with the fire duration of between 1 and 60 days. The "min" notation describes the desirability of the shortest elapsed time in the interval.

There is a large community of researchers and developers who have created algorithms for automatically creating, updating and managing temporal plans. Extensions to the core framework for dealing with preferences and uncontrollable events have also been developed. Rather than concentrate on the core planning capability, the remainder of our focus here will be on ways of integrating this core capability into a complete system for formulating and submitting science campaigns.

4. ARCHITECTURES FOR CAMPAIGN MANAGEMENT

The general vision here is of an *interface for coordination* (IFC) between a collection of observation resources (Earth observing sensors) and a community of potential users of one or more of those resources (viz. the Earth science community). Users of the IFC will formulate one or more campaign requests involving any of the resources in the collection. We call this organization *Multi-User Multi-Sensor* (MUMS). The collection of services provided by the IFC is somewhat similar to that of Earth science data archive retrieval systems such as EOSDIS. The primary difference, of course, is that the IFC would enable users to directly *plan* observation activities, rather than retrieve data that have been already acquired. The motivation for solving coordination of data acquisition at the planning phase is two-fold: more effective management of sensing resources through the simultaneous deliberative planning of all the resources together; and, from the perspective of the Earth scientist, the potential for higher utility data products through the ability to more effectively control what is observed, as well as when and how.

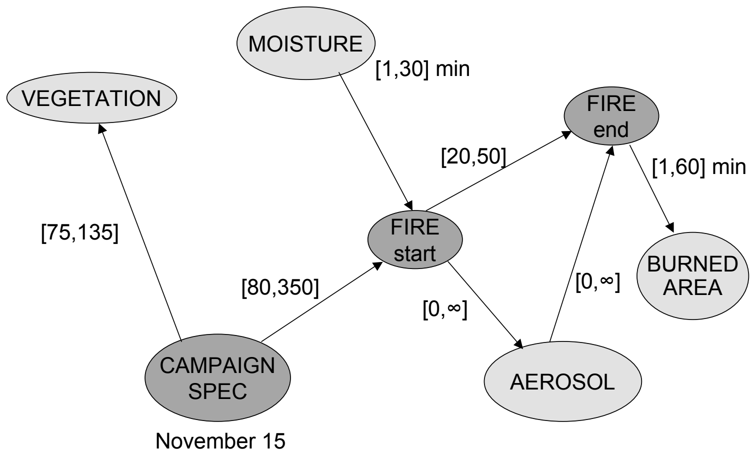

Figure 1. Diagram of a flexible plan for the fire scenario. Nodes are abstract versions of activities that are instantiated as specific observation descriptions (in light grey) or exogenous events (in dark grey).

Currently, observation planning tends to be Multi-User Single-Sensor.[12,13] Campaign planning involving multiple sensors is done, if at all, "manually," that is by individuals using multiple asset request interfaces, the telephone and email. For example, requests for NASA aircraft campaigns is done through a formal call to the community in which responses are gathered on a from that lists sensor type, platform, location, overflight time period, weather conditions, cloud cover, flight line orientation, sun angle limits, coincident satellite overpass, ground resolution, and tidal cycle.

There are a number of different architectural "styles" for a MUMS IFC, depending on the relationship between the IFC and individual missions. At one extreme is a completely centralized set of planner and scheduler services; this architecture is proposed in.[14] Here all planning for the collection of satellites is performed by a single system. The primary advantage of the centralized approach is that a single system has visibility into all the constraints on resources available to services a set of campaigns; this visibility means that a highly optimized schedule (with respect to resource utilization) can be produced. The primary disadvantage of this architecture is that missions lose control over science planning, which makes it highly unlikely that they would agree to the service provided by the interface.

This disadvantage suggests a more "culturally correct" architecture in which planning is distributed between the interface and the individual missions. An extreme instantiation of this architecture is one in which all control of planning is done by individual missions, and the interface plays the role of "bulletin board", posting campaign requests from users, hoping that missions will schedule them. Although more realistic from a "cultural" point of view, the bulletin board may provide too weak an interface between users and missions. Requesters may be more likely to use a system that is more tightly woven into the daily planning activities of missions.

The architecture we are developing provides a middle ground between the centralized and the purely distributed approach to planning. Adopting the core idea of the distributed approach, each sensor is managed by a separate mission as part of daily mission scheduling activity.[13] It is assumed that missions are fundamentally "uncooperative" in the sense that each does its science planning independently of the others, with little or no direct (peer to peer) coordination of activities. Consequently, the IFC imposes a hierarchical organization with respect to the missions. Further, it is assumed that individual missions are unwilling to relinquish control of the planning process for the instruments they manage; however, they are likely to accept a system that facilitates additional coordination by proposing incremental changes to their mission plans.

One way to implement this idea is to assume that each mission has a *plan database* of requests that provides the inputs to daily mission scheduling activity – such a database exists for Landsat 7, for example. One service provided by the IFC would be to generate a set of observation requests from the flexible campaign plans managed by the planner. Each observation request on a specific sensor would be translated into the format of the plan database of the missions managing the sensor. The IFC's "request manager" would facilitate an update to the plan database in time for the request to be considered as part of the mission daily scheduling activity. In this way, unlike the "bulletin board" model, the request could be treated on a par with other requests in the mission plan, thus likely improving its chances to be serviced. Nonetheless, the mission would maintain control of the daily scheduling activity, and the IFC would be seamlessly integrated into the community of resources.

In summary, the IFC would provide a collection of services that would enable campaign management as comprised of the following steps[4]:

1. The user specifies the set of observations and their constraints that comprise the campaign
2. The system generates and displays a *flexible temporal plan* based on this input; the user adds further constraints as desired, based on the information from the flexible plan.
3. The system displays a subset of possible fixed plans (sequences of observation requests) that are consistent with all the constraints specified, as requested by the user.
4. The user selects from among the list of fixed plans the one(s) that are most preferred.
5. The system proceeds to execute the requested fixed plan by incrementally updating mission plan databases.
6. The system notifies the user of the status of the requests, which may trigger additional changes to the campaign plan.

Notice that this solution requires some modeling of the resources managed by the individual missions at the IFC level. The knowledge contained in the model is required in order to ensure that the observation request is routed to the mission that is at least potentially able to service it. Minimally, to achieve this acceptable level of performance, the IFC must know sensor properties and orbit paths of the the satellites on which the sensors reside. More robust models could include information about the SSR (on-board storage), the dynamics of the device for slewing the sensor (for satellites in which the instrument can be independently slewed off nadir), and information about the duty cycle for the instrument. There are other variations on this basic IFC model. For example, some missions might not wish to grant an external system having direct access to their database. An alternative would be to simply submit a request to a mission planner point of contact by email. Finally, this architecture is potentially flexible enough to allow different services to be provided to different missions. For example, a mission like Landsat 7 that has robust scheduling tools may desire no scheduling services to be provided by the IFC, whereas another mission that has fewer automated planning services of its own may desire the IFC to do more planning for its own sensor.

5. RELATED WORK

The concept of considering the collection of Earth observation satellites jointly so as to exploit substitution possibilities and synergistic observations and allow scientists to gather data about multiple geophysical variables has been called the "sensorweb." Another sensorweb related idea is the use of data from one sensor to trigger the acquisition of data from another sensor. This is one way uncontrollable events can be handled. Global, coarse resolution sensors (such as MODIS) continually collect images of the surface. On-board or on the ground processing algorithms may be developed to find certain phenomena, such as cloud-free regions or active fires, whose locations can be used in geographic specifications to observation requests from finer resolution sensors. Several efforts are testing these ideas with NASA's Earth Observing-1, EO-1.[15]

6. CONCLUSIONS

Currently, scientists essentially carry out *post hoc* Earth observation campaigns by using satellite image archive retrieval systems such as the EOS Data Information System Data Gateway. Archive systems provide information on which biophysical or geophysical quantities have been collected in the past. By considering campaigns as a future planning problem and using a planning and scheduling system to make requests to satellite constellations, scientists may gain greater ability to study key Earth science questions.

7. ACKNOWLEDGEMENT

This research was supported in the AIST program by the Earth Science Technology Office.

REFERENCES

1. K. O. J.C. Baker and R. Williamson, eds., *Commercial Observation Satellites: At the Leading Edge of Global Transparency*, RAND and American Society of Photogrammetry and Remote Sensing, 2001.
2. G. on Earth Observations, "From observation to action: Achieving comprehensive, coordinated, and sustained earth observations for the benefit of humankind: Framework for a 10-year implementation plan," 2004. http://earthobservations.org.
3. P. Morris, N. Muscettola, and T. Vidal, "Dynamic control of plans with temporal uncertainty," in *Proceedings of the 17th International Joint Conference on A.I. (IJCAI-01)*, Morgan Kaufmann, San Francisco, CA, 2001.
4. R. Morris, J. Dungan, L. Khatib, J. Gasch, P. Hempel, J. W. illiams, and T. Wood, "Coordinated science campaign planning for earth observing missions," in *Earth Science Technology Conference*, NASA Earth Science Technology Office, Palo Alto, CA, 2004.
5. G. Zhou, O. Baysal, J. Kaye, S. Habib, and C. Wang, "Concept design of future intelligent earth observing satellites," *International Journal of Remote Sensing* **25**, pp. 2667–2685, 2004.
6. B. Huebert, T. Bates, P. Russell, G. Shi, Y. Kim, K. Kawamura, G. Carmichael, and T. Nakajima, "An overview of ACE-Asia: Strategies for quantifying the relationships between Asian aerosols and their climatic impacts," *Journal of Geophysical Research* **108**, p. doi:10.1029/2003JD003550, 2003.
7. M. King, ed., *EOS Science Plan: The State of Science in the EOS Program*, NASA Goddard Space Flight Center, Greenbelt, Maryland, 1999.
8. "Sourcebook," in *Aviation Week*, 2004.
9. S. M. R.A. Ryerson and A. Budge, eds., *The Manual of Remote Sensing: Earth Observing Platforms & Sensors*, American Society of Photogrammetry and Remote Sensing, 1997. 3rd edition, CD-ROM.
10. W. A. Stoney, *Guide to land imaging satellites*, American Society of Photogrammetry and Remote Sensing, 2004. http://www.asprs.org/asprs/news/satellites/satellites.html.
11. J. S. Reid, E. M. Prins, D. L. Westphal, S. Christopher, C. Schmidt, K. A. Richardson, M. Theisen, E. A. Reid, and T. Eck, "Flambe: The fire locating and modeling of burning emissions project," in *Proceedings of the Fall Meeting of the American Geophysical Union*, San Francisco, CA, 2001.
12. H. Muraoka, R. Cohen, T. Ohno, and N. Doi, "ASTER observation scheduling algorithm," in *Proceedings of Space Ops 98*, Tokyo, Japan, 1998.
13. W. Potter and J. Gasch, "A photo album of earth: Scheduling daily Landsat 7 activities," in *Proceedings of Space Ops 98*, Tokyo, Japan, 1998.
14. J. Dungan, J. Frank, A. Jónsson, R. Morris, and D. Smith, "Advances in planning and scheduling of remote sensing instruments for fleets of earth orbiting satellites," in *Earth Science Technology Conference*, NASA Earth Science Technology Office, Pasadena, CA, 2002.
15. S. Chien, "Autonomous science on the EO-1 mission," in *Proceedings of International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, Nara, Japan, 2003.